arXiv:1612.04689v1 [cs.DS] 14 Dec 2016

An Integer Interior Point Method for Min-Cost Flow Using Arc Contractions and Deletions

Ruben Becker[1,2,3] and Andreas Karrenbauer[1,2] and Kurt Mehlhorn [1]

firstname.lastname@mpi-inf.mpg.de

[1] Max Planck Institute for Informatics, Saarbrücken, Germany

[2] Max Planck Center for Visual Computing and Communication

[3] Saarbrücken Graduate School for Computer Science

Abstract

We present an interior point method for the min-cost flow problem that uses arc contractions and deletions to steer clear from the boundary of the polytope when path-following methods come too close. We obtain a randomized algorithm running in expected $\tilde{O}(m^{3/2})$ time that only visits integer lattice points in the vicinity of the central path of the polytope. This enables us to use integer arithmetic like classical combinatorial algorithms typically do. We provide explicit bounds on the size of the numbers that appear during all computations. By presenting an integer arithmetic interior point algorithm we avoid the tediousness of floating point error analysis and achieve a method that is guaranteed to be free of any numerical issues. We thereby eliminate one of the drawbacks of numerical methods in contrast to combinatorial min-cost flow algorithms that still yield the most efficient implementations in practice, despite their inferior worst-case time complexity.

1 Introduction

The min-cost flow problem is one of the most general network flow problems that can be solved in polynomial time. It generalizes the max-flow as well as the shortest path problem. Moreover, it has many applications, since a lot of real-world problems can be modeled as a min-cost flow problem. Non-surprisingly there is a rich literature concerning the problem. The work can roughly be split into two major lines. First, there is the line of combinatorial algorithms that starts with the seminal work of Ford and Fulkerson [11] and continues with Edmonds and Karp [10], who gave the first polynomial time algorithm for the problem. Further polynomial time algorithms are the strongly polynomial time algorithm of Orlin [24] and the scaling approaches [16, 15] that peak in the double scaling technique by Ahuja et al. [2] achieving a bound of $O(nm \log\log U \cdot \log(nC))$, where n denotes the number of nodes, m the number of arcs, U the maximal capacity, and C the maximal cost in the network. The second line of work originates in the results of Vaidya [28] who was the first to specialize interior point methods to network-flow problems. Lagging behind the combinatorial methods at that time, this line of research has become extremely fruitful in the last decade; most prominently with the work of Daitch and Spielman [9] who gave a randomized $\tilde{O}(m^{3/2})$-time algorithm.[1] They were the first to use nearly-linear time equation solvers [26, 20, 17] in the context of generalized min-cost flow problems for approximately solving systems of linear equations in each of $\tilde{O}(\sqrt{m})$ iterations, thereby, for the first time, achieving sub-quadratic bounds in sparse graphs and under the similarity assumption, i.e., when the input data is polynomially bounded in n. The same running time (up to log factors) was achieved in [6] for the classical min-cost flow problem with a simple potential reduction algorithm. Even more recently, the breakthrough result of Lee and Sidford [22] for solving general linear programs in rank-many

[1] The tilde indicates that logarithmic factors in n, U, C are hidden.

iterations yields a very intricate, randomized, interior point method of complexity $\tilde{O}(\sqrt{n}m)$ for the min-cost flow problem. Most recently, Cohen et al. developed a new framework to analyze interior point methods [8]. They obtain a bound of $\tilde{O}(m^{3/2})$ for the case of min-cost flow problems with the restriction that there are only unit capacities.

Although achieving the currently best-known bounds in theory, the interior point methods lack behind the combinatorial methods in practice, see for example [13, 12, 5] for efficient implementations of combinatorial algorithms and [19] for an extensive comparison of implementations. We are not aware of any interior point method that was reported to be competitive with the combinatorial approaches on real world problems. The reasons are many. First, combinatorial algorithms are rather easy to implement and necessary subroutines are readily available in standard libraries, whereas the ingredients that are required to achieve the better theoretical results are rather involved and use complicated subroutines such as low-stretch spanning trees [1] or spectral vertex sparsifiers [21] that are already difficult to implement efficiently on their own. Second, interior point methods are formulated for real arithmetic and implementations require high-precision floating point numbers in order to avoid numerical instabilities, even if all input data is integer. In contrast to this, the combinatorial algorithms have the advantage that, given an integral input, they get along with fixed-precision integer arithmetic, which makes them very robust in practice. Note that, if floating-point-arithmetic is used, even combinatorial algorithms may suffer from numerical instability as was observed by Althaus and Mehlhorn in LEDA [4]. In this paper, we address the second issue of practical implementations of interior point methods decribed above by giving an interior point method that uses integer arithmetic on numbers of polynomial length. We proceed by defining the min-cost flow problem formally.

An instance of the min-cost flow problem is defined by a directed graph $G = (V, A)$ with $|V| = n$ and $|A| = m$, node demands $b \in \mathbb{Z}^n$ with $\mathbb{1}^T b = 0$, arc costs $c \in \mathbb{Z}^m$ and arc capacities $u \in \mathbb{N}^m$. We define $\beta := \gcd(b, u)$ and $\gamma := \gcd(c)$, i.e. greatest common divisors of the entries of the respective vectors. Moreover, we use $U := \max\{\|u\|_\infty/\beta, \|b\|_1/(2\beta)\}$ and $C := \|c\|_\infty/\gamma$. A flow $x \in \mathbb{R}^m$ is called *feasible*, if $0 \leq x \leq u$ and $x(\delta^{\text{in}}(v)) - x(\delta^{\text{out}}(v)) = b_v$ for every $v \in V$. [2] The task is to find the feasible flow x^* of minimal cost, i.e., $c^T x^* \leq c^T x$ for all feasible flows x, or to assert that no feasible flow exists. It is well-known, see, e.g., [3, 6], that we can restrict ourselves to the problem being feasible, having only non-negative costs, and no capacity constraints. We will also justify the last assumption in Appendix A.

Assuming the problem to be uncapacitated, we will use the notation (G, b, c) for a min-cost flow instance on the graph G, with demand vector b and costs c, which can be written as a primal-dual linear programming pair

$$\min\{c^T x \ : \ Ax = b \text{ and } x \geq 0\} = \max\{b^T y \ : \ A^T y + s = c \text{ and } s \geq 0\}, \tag{1}$$

where $A \in \{-1, 0, 1\}^{n \times m}$ denotes the node-arc incidence matrix of G, i.e., A contains a column α for every arc $(v, w) \in A$ with $\alpha_v = -1$, $\alpha_w = 1$ and $\alpha_i = 0$ for all $i \notin \{v, w\}$.

Path following methods maintain a pair of primal/dual feasible solutions x and (y, s) such that $x_a s_a \approx \mu$ for all $a \in A$ for a positive parameter μ, which they drive to zero. They thereby follow the so-called *central path*, which is the set of all points satisfying $x_a s_a = \mu$ for all $a \in A$. As a consequence of μ approaching zero, the *duality gap* $c^T x - b^T y = x^T s \approx \mu m$ goes to zero. It may however happen, that either x_a or s_a becomes tiny while the other becomes huge. Such behavior is numerically unpleasant as it requires to simultaneously deal with huge and tiny numbers.

1.1 Our Contribution

In this paper, we resolve the issue of simultaneously having to deal with huge and tiny numbers in interior point methods for min-cost flow by only considering the arcs with sufficiently large values of x_a and s_a. That is, whenever x_a becomes too small, we delete the arc a from the network and whenever s_a becomes too small, we contract the arc a. We thereby obtain a reduced

[2] We write $f(S) := \sum_{a \in S} f_a$ for $S \subseteq A$ for any vector $f \in \mathbb{R}^m$ and for a node $v \in V$, we use $\delta^{\text{in}}(v)$ for the ingoing arcs of v and $\delta^{\text{out}}(v)$ for its outgoing arcs, respectively.

problem in form of a minor of the original network. We do not revoke any of these operations until termination of our algorithm and show that a near-optimal solution for the reduced problem can be lifted efficiently to an optimal solution of the original network. Arc deletion and contraction has been used in combinatorial optimization before [27, 25], but, to our knowledge, not in combination with path-following methods for network flow problems.

For the reduced problem, we can show that one can stay sufficiently close to the central path by hopping over points on an integer lattice in the vicinity of the central path. More precisely, we show that if the greatest common divisor (gcd) of the demands and capacities and the gcd of the costs are sufficiently large, we can perform all arithmetic operations on integers as typical for the practically superior combinatorial algorithms. The requirement on the gcds can be easily achieved by an initial scaling of the demands, capacities and costs. Such scaling comes at a very low price because the expected running time of $\tilde{O}(m^{3/2})$ only depends logarithmically on these scaling factors. We remark that it appears obvious that integer arithmetic can be achieved by scaling with a huge number that is, say, doubly exponential in m, n. However, we show that scaling with numbers that are polynomial in n is sufficient. Moreover, we manage to make the exact numerical requirements of our algorithm explicit by stating absolute bounds on the appearing numbers, without hiding anything in O-notation.

Theorem 1. *Optimal primal and dual solutions for a min-cost flow problem can be computed in expected $\tilde{O}(m^{3/2})$ time using only integer arithmetic on numbers bounded by $2^{31}m^{10}U^2C^2$, where U and C are bounds on the maximum capacity and cost, respectively.*

Note that the running time of our algorithm matches the time bound of Daitch and Spielman [9], which dominates the known bounds for combinatorial algorithms in many settings. For a better comparison, we remark that Daitch and Spielman [9] give a bound on the bit length of the numbers appearing in their algorithms, as well. They state that the methods work on numbers of bit length at most $O(\log(nU/\varepsilon))$, where $\varepsilon = (12m^2U^3)^{-1}$ for solving the standard min-cost flow problem exactly [3]. If the hidden constant in the O-notation is known, say, to be k, then their method could also be used with integer arithmetic on numbers bounded by $12^k n^k m^{2k} U^{4k}$. For robust implementations, it is advantageous to know the exact precision requirements, as an appropriate scaling can be performed in the beginning without having to worry about numerical issues or overflows in the course of the algorithm. Alternatively, one can still use adaptive precision, check the correctness of the output, and repeat the computation with higher precision, if necessary. However, this approach usually requires changes to the implementation as insufficient precision may lead to diverging behavior of sub-routines, see for example [18]. Adding a timer is no solution either as it turns diverging behavior into worst-case behavior. However, such precautions are mandatory and the only choice, if the bounds are not known explicitly.

For the potential-reduction method from [6], we only gave a bound on the number of arithmetic operations and a strict bound on the size of the emerging numbers is non-trivial to obtain, if not even impossible. It is essential for our proof in this paper to exploit graph minors and to use a path-following method instead. Furthermore, it is interesting to note that our lower bound on the value of primal variables does not depend on the costs, capacities, or demands in the network, but only on n and m. This might be interesting in the context of strongly polynomial time algorithms for the capacitated min-cost flow problem, where the current record bound is $\tilde{O}(m^2)$, due to Orlin [24].

We want to remark that we consider this work rather as a first important step in the direction of making numerical requirements of interior point methods for combinatorial problems apparent, than as giving the ultimate final answer to practitioners. We feel that our paper, despite the fact that it treats all numerical details, is still fairly elegant. This is due to the fact that we compute in integers, and by this, avoid the tediousness of floating point error analysis.

[3] Note that in the case of generalized min-cost flow, the algorithms of Daitch and Spielman only give ε-approximate solutions in general and thus the numerical requirements will also depend on the desired ε. For the standard min-cost flow problem they show that an ε of the size as described above is sufficient.

Overview At the beginning of Section 2, we describe the high-level idea of the interior point method and explain how numerical issues come into play. In Subsection 2.1, we present our way of dealing with this problem, deletion and contraction of arcs, and treat some theoretical observations that allow us to pursue this way. We end this section with the concept of a so-called proxy. We then describe our algorithm in more detail and also give a pseudocode implementation of it, see Subsection 2.2. The interior points that the algorithm maintains are actually interior points of a slightly perturbed instance, we describe this in Subsection 2.3, where we will see that an optimal tree solution of the perturbed instance will be sufficient for obtaining an optimal tree solution for the original problem. Subsection 2.4 describes how to obtain this tree solution, namely using the crossover algorithm from [6, Section 2], see also [5] for more details on the crossover algorithm. We treat the centering step and the related numerical issues in Section 3. Finally, we conclude in Section 4.

2 Integer Interior Point Algorithm

Our method belongs to the class of path-following interior point methods. These methods maintain a pair of primal/dual feasible solutions $x, s \in \mathbb{R}^m$ while more or less uniformly approaching the complementary slackness conditions. That is, the conditions are parameterized uniformly by a positive μ, i.e., $x_a s_a = \mu$ for all $a \in A$, s.t. one obtains a pair of primal/dual optimal solutions in the limit $\mu \to 0$. It can be shown [29] that the points satisfying theses conditions form a curve called *central path*, which is contained in the interior of the primal/dual polyhedron. Moreover, moving along the central path with decreasing μ concurrently decreases the *duality gap* $c^T x - b^T y = x^T s = \mu m$.

As said above, path-following methods typically only maintain interior points that are sufficiently close to the central path. In this paper, we measure the closeness to the central path w.r.t. some $\mu > 0$ by the relative deviation in the 1-norm $||\sigma||_1$ where $\sigma \in \mathbb{R}^m$ and $\sigma_a := \frac{x_a s_a}{\mu} - 1$. A sufficiently small deviation, e.g., $||\sigma||_1 \leq \delta$ for some constant $\delta < 1$, guarantees that both x_a and s_a remain positive because $||\sigma||_1 \leq \delta$ implies that $x_a s_a / \mu \geq 1 - \delta$ is well separated from 0. However, this does not prohibit that either x_a or s_a becomes tiny while the other becomes huge. Such behavior is not only numerically unpleasant, it is also an obstacle for using fixed-precision integer arithmetic.

2.1 Minors to the Rescue

We deal with the issue mentioned above in a somewhat radical way: We only consider the arcs with sufficiently large values for x_a and s_a. That is, whenever x_a becomes too small, we delete the arc from the network and whenever s_a becomes too small, we contract the arc a. We thereby obtain a minor of the original network. This is captured formally by the following definition.

Definition 1. *Let $G = (V, A)$ be a directed graph with m arcs. For $\varepsilon > 0$ and $x, s \in \mathbb{R}^m_{\geq 0}$, we call the minor H of G obtained by deleting the arcs in $D := \{a \in A : x_a < \frac{\varepsilon}{m}\}$ and contracting the arcs in $C := \{a \in A : s_a < \frac{\varepsilon}{m}\}$ the ε-x-s-minor of G.*

This approach might be surprising, but it has its foundation in classical min-cost flow theory [3, Section 10.6]. We provide a brief reasoning for being self-contained. We first review a well-known useful fact from general LP duality.

Lemma 1. *Let $\min\{c^T x : Ax = b, x \geq 0\}$ be bounded and feasible and let $\max\{b^T y : A^T y + s = c, s \geq 0\}$ denote its dual linear program.*

- *A primal optimal solution x^* is also an optimal solution of $\min\{\bar{s}^T x : Ax = b, x \geq 0\}$ for any dual feasible solution $\bar{y}, \bar{s}$.*
- *A dual optimal solution y^*, s^* is also an optimal solution of $\min\{\bar{x}^T s : A^T y + s = c, s \geq 0\}$ for any primal feasible solution $\bar{x}$.*

Proof. Let x^* and y^*, s^* be primal and dual feasible solutions, respectively.

- Let $\bar{y}, \bar{s}$ be a dual feasible solution. Define $y' := y^* - \bar{y}$ and $s' := s^* = c - A^T y^*$, then $A^T y' + s' = A^T y^* - A^T \bar{y} + s^* = c - A^T \bar{y} = \bar{s}$ and hence y', s' are feasible for $\max\{b^T y : A^T y + s = \bar{s}, s \geq 0\}$. Moreover, it holds that $x^{*T} s' = x^{*T} s^* = 0$. It follows that x^*, y', s' are primal and dual optimal for $\min\{s^T x' : Ax' = b, x' \geq 0\}$ and $\max\{b^T y' : A^T y' + s' = s, s' \geq 0\}$, respectively. This shows the claim.

- Let $\bar{x}$ be a primal feasible solution, i.e., $A\bar{x} = b$ holds. Suppose for contradiction that

$$\begin{aligned}\bar{x}^T s^* &> \min\{\bar{x}^T s : A^T y + s = c, s \geq 0\} = \min\{\bar{x}^T (c - A^T y) : A^T y + s = c, s \geq 0\} \\ &= c^T \bar{x} - \max\{b^T y : A^T y + s = c, s \geq 0\} = c^T \bar{x} - b^T y^* = \bar{x}^T s^*.\end{aligned}$$

□

The following lemma shows the rationale for deleting arcs with tiny flow x and thus huge reduced costs s in the vicinity of the central path (respectively, contracting arcs with tiny reduced costs and huge flow). [4]

Lemma 2. *Let $x, s \in \mathbb{R}^m$ be primal/dual feasible solutions of (G, b, c) with $b \in \mathbb{Z}^n$ and $c \in \mathbb{Z}^m$.*

- *If $s_a > x^T s$ for some $a \in A$, then $x_a^* = 0$ for every optimal solution x^*.*
- *If $x_a > x^T s$ for some $a \in A$, then $s_a^* = 0$ for every optimal solution s^*.*

Proof. Because of the total unimodularity of the node-arc incidence matrix A all basic solutions are integral. Thus, it suffices to show the claims for optimal integral basic solutions. All other optimal solutions are convex combinations of these.

- Assume for contradiction that $x_a^* > 0$ in an optimal integral solution, i.e., $x_a^* \geq 1$. Then $x^{*T} s \geq x_a^* s_a > x^T s$, this shows that x^* is not optimal for $\min\{s^T x : Ax = b, x \geq 0\}$, contradicting Lemma 1.

- Assume for contradiction that $s_a^* > 0$ in an optimal integral solution, i.e., $s_a^* \geq 1$. Then $x^T s^* \geq x_a s_a^* > x^T s$, this shows that s^* is not optimal for $\min\{x^T s : A^T y + s = c, s \geq 0\}$, contradicting Lemma 1. □

Note that the central path condition $\sum_{a \in A} |\frac{x_a s_a}{\mu} - 1| \leq \delta$ implies $x_a s_a \in [(1-\delta)\mu, (1+\delta)\mu]$ for any arc $a \in A$. Now assume that for some arc $a \in A$, the value of x_a becomes tiny, i.e., $x_a < \frac{\varepsilon}{m}$. Then $s_a \geq \frac{(1-\delta)\mu}{x_a} > \frac{(1-\delta)\mu m}{\varepsilon} \geq (1+\delta)\mu m \geq x^T s$ follows for $\varepsilon := \frac{1-\delta}{1+\delta} < 1$. Thus Lemma 2 applies and we can conclude that $x_a^* = 0$ in every optimum primal solution x^*. A completely symmetric approach shows that, if $s_a < \frac{\varepsilon}{m}$ for an arc $a \in A$, then $s_a^* = 0$ in every optimal dual solution s^*.

The analyses of typical classical combinatorial algorithms that use arc deletion or arc contraction, respectively, are based on an argument that termination is guaranteed because at least one arc has to deleted or contracted after a certain amount of iterations and there are only m arcs. However, we have a different termination criterion that mainly depends on the duality gap. In [6], we showed that it suffices to compute a pair of primal/dual interior points with duality gap strictly less than 1 to efficiently perform a crossover to an optimum integral basic feasible solution. In this paper, we extend this result and show that proximity to the optimum solution w.r.t. the minor suffices. To this end, we first formally define what proximity means in our context.

Definition 2. *A pair of primal and dual feasible solutions x, s of (G, b, c) with $b \in \mathbb{Z}^n$ and $c \in \mathbb{Z}^m$ is called a proxy for the optimum of (G, b, c) if $\sum_{a \in A_H} x_a s_a < (1-\varepsilon)^2$ where A_H is the arc-set of an ε-x-s-minor of G for some $\varepsilon < 1$.*

[4] Note that this result can be extended to general integral linear programs, but we prefer to keep it simple for the sake of presentation.

2.2 The Algorithm

In Appendix A, we show how to construct, in linear time, initial interior points for an auxiliary instance with the same optimum that fulfill the central path condition $\|\sigma\|_1 \leq \delta$ for $\mu_0 = 3m\beta\gamma UC/\delta$ and thus can be used as input for Algorithm 1. The idea is to decrease the current μ by roughly a factor of $1-\delta/\sqrt{m}$, more precisely, we set $\mu \leftarrow \lceil(1-\tau)\mu\rceil$, where $\tau = \delta/\sqrt{m}$, as long as we have not achieved a solution that is a proxy for the optimum of (G, b, c) yet. Then, we restore closeness to the central path w.r.t. the new μ and the minor in a so-called *centering step*. Observe that the centering step and the absence of tiny x_a and s_a in the minor implies that their respective counterparts are bounded by $(1+\delta)\mu/\varepsilon$. Hence, the magnitude of the initial μ yields an upper bound on the numbers that have to be considered in this algorithm.

Moreover, let $\beta := \gcd(b)$ and $\gamma := \gcd(c)$. Observe that the instance $(G, b/\beta, c/\gamma)$ still has integral demands and costs and it is combinatorially equivalent to the instance (G, b, c). However, we do not scale the instance down, but we rather maintain this granularity to restrict the computations to integer arithmetic, i.e., we follow the central path by hopping from one point of the integer grid to an other one. We show that this is possible for sufficiently large β and γ and that the running time only scales logarithmically in β and γ so that up-scaling to meet this condition can be achieved at low cost. More precisely, suppose that a given input instance does not satisfy our assumption of having sufficiently large gcd's. Then, the approach is to scale down the instance by dividing by $\gcd(b)$ and $\gcd(c)$, respectively, and scaling them up by sufficiently large numbers that are given by the following theorem.

Theorem 2. *Let (G, b, c) be a min-cost flow instance with $b \in \beta\mathbb{Z}^n$ and $c \in \gamma\mathbb{Z}^m$. There is a randomized algorithm to compute $x, s \in \mathbb{Z}^m$ s.t. $x/\beta, s/\gamma$ is a proxy for the optimum of $(G, b/\beta, c/\gamma)$ in expected $\tilde{O}(m^{3/2})$ time using only integer arithmetic on numbers bounded by $2^8 m^3 \gamma UC$, provided that $\beta \geq 2^8 m^3$ and $\gamma \geq 2^{15} m^4 \beta UC$.*

Algorithm 1: Integer Interior Point Algorithm

Input : (G, b, c) with $\gcd(b) = \beta$ and $\gcd(c) = \gamma$, feasible interior $x, s \in \mathbb{Z}^m$, and $\mu > 0$ s.t. $\|\sigma\|_1 := \sum_{a\in A} |\frac{x_a s_a}{\mu} - 1| \leq \delta$.

Output : Vectors x, s s.t. x/β and s/γ yield a proxy for the optimum of $(G, b/\beta, c/\gamma)$.

repeat
- $H = \texttt{minor}(G, x, s)$
- $\mu \leftarrow \lceil(1-\tau)\mu\rceil$, where $\tau = \delta/\sqrt{m}$
- $x, s = \texttt{centering_step}(H, x, s, \mu)$

until $\sum_{a\in A_H} x_a s_a < (1-\varepsilon)^2\beta\gamma$

return x, s

Algorithm 2: $\texttt{minor}(G, x, s)$

Let $\varepsilon := \frac{1-\delta}{1+\delta}$.

for $a \in A$ **do**
- **if** $x_a < \frac{\varepsilon\beta}{m}$ **then** remove a from G **if** $s_a < \frac{\varepsilon\gamma}{m}$ **then** contract a in G

return G

We sketch the algorithm that achieves the above theorem as pseudo-code in Algorithms 1 and 2. The randomization only happens in the subroutine for the centering step, where fundamental cycles w.r.t. a low-stretch spanning tree are sampled. We discuss the centering step in detail in Section 3. For now it is only important that it maintains closeness to the central path w.r.t. the arcs in the minor. For the sake of presentation, we did not maintain the y-variables that can be used to restore the reduced costs on the arcs that have been deleted. But this can be easily achieved. Moreover, one can extend the cycles mentioned above from the minor through the contracted nodes to cycles in the original graph such that one can also keep track of the updates to maintain a feasible solution in the original graph. We thereby can obtain a proxy for the optimum of (G, b, c). However, the variables of the arcs in the minor describe an interior point of a slightly perturbed instance.

2.3 Perturbation

Suppose, for example, that we have an ε-x-s-minor H that differs from G just by the deletion of a single arc $\hat{a}$. Hence, the set of nodes of H coincides with the one of G. However, projecting x to x_H, s to s_H, and c to c_H by removing the entry corresponding to $\hat{a}$ does not yield a pair of primal/dual feasible solutions x_H, s_H for (H, b, c_H), but it is rather a feasible solution to $(H, \tilde{b}, c_H)$ where $\tilde{b} = b - Ae_{\hat{a}}e_{\hat{a}}^T x$ and we denote with $e_{\hat{a}}$ the vector of unity corresponding to the index $\hat{a}$. Note that the two demands vectors b and $\tilde{b}$ are almost identical because $x_{\hat{a}}$ is tiny. We first define formally what we mean by almost identical instances and then draw the connection to minors.

Definition 3. *An instance $(G, \tilde{b}, \tilde{c})$ is called* ε-perturbed *instance of (G, b, c), if $\|b - \tilde{b}\|_1 \leq 2\varepsilon$ and $\|c - \tilde{c}\|_1 \leq \varepsilon$.*

We introduce the following notation to transform instances and solutions:

- Denote with $I_b \in \mathbb{R}^{n_H \times n}$ the linear map that transforms b into b_H, where for every node $v \in V_H$, $b_H(v) = \sum_{u \in C_v} b(u)$ for $C_v \subseteq V$ being all nodes in G that get contracted to u.
- As above for any $a \in A$, we denote with $e_a \in \{0,1\}^m$ the vector of unity that has a 1 in its a'th component and 0 elsewhere, where the arcs are ordered as the columns of the node-arc-incidence matrix A. Similarly, with $e_v \in \{0,1\}^n$ we mean the vector of unity that has a 1 in its v'th component, 0 elsewhere and the considered ordering is as in the rows of A.
- Denote as $M_H \in \{0,1\}^{m_H \times m}$ the map with the a'th row being e_a for all $a \in A_H$.
- Denote as $N_H \in \{0,1\}^{n_H \times n}$ the map with the v'th row being e_v for all $v \in V_H$.
- For some $S \subseteq A$, let $I_S := \sum_{a \in S} e_a e_a^T$.

Lemma 3. *Let x, s be a pair of primal/dual feasible solutions of (G, b, c) let $C := \{a \in A_G : s_a < \frac{\varepsilon}{m}\}$ and $D := \{a \in A_G : x_a < \frac{\varepsilon}{m}\}$ for some $0 < \varepsilon < 1$. Then $(G, \tilde{b}, \tilde{c})$ with $\tilde{b} = b - AI_D x$ and $\tilde{c} = c - I_C s$ is an ϵ-perturbed instance of (G, b, c).*

Proof. The proof is a simple calculation:

$$\|b - \tilde{b}\|_1 = \|AI_D x\|_1 \leq 2\|I_D x\|_1 \leq 2\varepsilon \quad \text{and} \quad \|c - \tilde{c}\|_1 = \|I_C c\|_1 \leq \varepsilon. \qquad \square$$

The following Theorem shows the equivalence of the original and the perturbed problem. For a min-cost flow instance (G, b, c) and a given spanning tree T of G, we call $x \in \mathbb{R}^m$ a *primal tree solution*, if $Ax = b$ and $x_a = 0$ for all $a \in A \setminus T$ and we call $y, s \in \mathbb{R}^{n+m}$ a *dual tree solution* if $A^T y + s = c$ and $s_a = 0$ for all $a \in T$.

Theorem 3. *Let $(G, \tilde{b}, \tilde{c})$ be an ε-perturbed instance of (G, b, c) for some $0 < \varepsilon < 1$ and let T be a spanning tree of G.*

a) If the unique tree solution w.r.t. T in $(G, \tilde{b}, \tilde{c})$ is primal feasible, then the unique primal tree solution w.r.t. T in (G, b, c) is feasible.

b) If the unique tree solution w.r.t. T in $(G, \tilde{b}, \tilde{c})$ is dual feasible, then the unique dual [5] tree solution w.r.t. T in (G, b, c) is feasible.

c) If the spanning tree T is optimal for $(G, \tilde{b}, \tilde{c})$, then it is optimal for (G, b, c).

Proof. a) Let x_T and $\tilde{x}_T$ denote the unique primal tree solutions corresponding to T in (G, b, c) and $(G, \tilde{b}, \tilde{c})$, respectively, i.e., $A_T x_T = b$ and $A_T \tilde{x}_T = \tilde{b}$ holds. Similarly let d be the unique primal tree solution corresponding to T for the right hand side $b - \tilde{b}$, i.e., $A_T d = b - \tilde{b}$. Note that in order to show feasibility of x_T, it suffices to show that $x_T \geq 0$. It follows that

$$A_T d = b - \tilde{b} = A_T(x_T - \tilde{x}_T)$$

[5] Here *unique dual* refers only to the reduced costs because $A^T(y + t\mathbb{1}) = A^T y$ for all $t \in \mathbb{R}$.

and thus by the uniqueness of d, we conclude $d = x_T - \tilde{x}_T$. It follows that $\|x_T - \tilde{x}_T\|_\infty = \|d\|_\infty \leq \|b - \tilde{b}\|_1/2 \leq \varepsilon$, using that $(G, \tilde{b}, \tilde{c})$ is an ε-perturbed instance. However, since $b \in \mathbb{Z}^n$ also $x_T \in \mathbb{Z}^m$ and since $\varepsilon < 1$ it follows that $\tilde{x}_T \geq 0$ implies $x_T \geq 0$.

b) Let y_T, s_T and $\tilde{y}_T, \tilde{s}_T$ denote the unique dual tree solutions corresponding to T in (G, b, c) and $(G, \tilde{b}, \tilde{c})$, respectively, i.e., $A^T y_T + s_T = c$, $s_T(a) = 0$ for $a \in T$ and $A^T \tilde{y}_T + \tilde{s}_T = \tilde{c}$, $\tilde{s}_T(a) = 0$ for $a \in T$ holds. Similarly let p, d be the unique dual tree solution corresponding to T for the right hand side $c - \tilde{c}$, i.e., $A^T p + d = c - \tilde{c}$. Note that in order to show feasibility of y_T, s_T, it suffices to show that $s_T \geq 0$. It follows that

$$A^T p + d = c - \tilde{c} = A^T(y_T - \tilde{y}_T) + s_T - \tilde{s}_T$$

and thus by the uniqueness of d, we conclude $d = s_T - \tilde{s}_T$. Furthermore for $a \in T$, we have $d(a) = 0$ and for $a \in A \setminus T$, it holds that $|d_a| = |c_a - \tilde{c}_a - (p_w - p_v)| = |c_a - \tilde{c}_a - \sum_{b \in P(v,w)} c_b - \tilde{c}_b| \leq \|c - \tilde{c}\|_1$. It follows that $\|s_T - \tilde{s}_T\|_\infty = \|d\|_\infty \leq \|c - \tilde{c}\|_1 \leq \varepsilon$, using that $(G, \tilde{b}, \tilde{c})$ is an ε-perturbed instance. However, since $c \in \mathbb{Z}^m$ also $s_T \in \mathbb{Z}^m$ and since $\varepsilon < 1$ it follows that $\tilde{s}_T \geq 0$ implies $s_T \geq 0$.

c) Assume the tree T is optimal for $(G, \tilde{b}, \tilde{c})$, i.e., the tree is primal and dual feasible (complementary slackness). Together with part a) and b), we conclude that T is also optimal for (G, b, c). □

We note that in the non-degenerate case, the reverse statements of a) and b) hold as well. As our algorithm will output a tree solution to the perturbed problem, the above Theorem is sufficient to show that the output tree is also an optimal tree for the original problem.

2.4 Crossover

The crossover algorithm from [6, Section 2] takes a min-cost flow instance $(G, \tilde{b}, \tilde{c})$ and primal/dual feasible points $\tilde{x}, (\tilde{y}, \tilde{s})$. Given that $\tilde{c}$ is integral, it outputs an integral dual tree solution y^*, s^* with $b^T y^* \geq b^T \tilde{y}$. In a nutshell, the crossover algorithm works as follows. It iteratively constructs nested cuts $\{s\} = S^1 \subset S^2 \subset \ldots \subset S^n = V$, starting with an arbitrary node s. During iteration i, the algorithm modifies the potentials y along the cut S^i ensuring that $b^T y$ does not decrease. This is done by incrementing the potentials of all nodes in S^i, if $b(S^i) \geq 0$, and by decrementing them otherwise. The potentials are increased (decreased) until the reduced cost of at least one arc a on the cut $S, V \setminus S$ becomes 0. Then, the node adjacent to S^i through a is being added to S^i. We note that the amount by which the potentials inside of S^i are changed relatively to the nodes outside of S^i, is exactly equal to the reduced cost of the arc a.

Theorem 4. *Let (G, b, c) be a min-cost flow instance with $b \in \beta\mathbb{Z}^n$ and $c \in \gamma\mathbb{Z}^m$. Given $x, s \in \mathbb{Z}^m$ s.t. $x/\beta, s/\gamma$ is a proxy for the optimum of $(G, b/\beta, c/\gamma)$, a pair of primal-dual optimum solutions x^*, s^* of (G, b, c) can be computed in $O(m^{3/2} \log \frac{n^2}{m} \log U)$ time.*

Proof. Let $\bar{x} = x/\beta$, $\bar{s} = s/\gamma$ as well as $\bar{b} = b/\beta$ and $\bar{c} = c/\gamma$. Consider the projection $\tilde{x}$ of $\bar{x}$ to the subspace with $\tilde{x}_a = 0$ for all $a \in D$ and the projection $\tilde{s}$ of $\bar{s}$ to the subspace with $\tilde{s}_a = 0$ for all $a \in C$. Note that these projections are primal/dual feasible for the ε-perturbation $(G, \tilde{b}, \tilde{c})$ of (G, b, c), where $\tilde{b} = \bar{b} - AI_D x$ and $\tilde{c} = \bar{c} - I_C s$ and their duality gap satisfies $\tilde{x}^T \tilde{s} = \sum_{a \in A_H} \bar{x}_a \bar{s}_a < (1 - \varepsilon)^2$. Given feasible $\bar{s}$, we can compute corresponding feasible $\bar{y}$ in linear time by propagation from the root to the leaves of an arbitrary spanning tree. The same holds for $\tilde{y}$ and $\tilde{s}$. If we perform the crossover procedure with $\tilde{y}, \tilde{s}$ in $(G, \tilde{b}, \tilde{c})$, we obtain a tree solution w.r.t. some spanning tree T. By Theorem 3, this spanning tree also yields a dual feasible solution for (G, b, c). Moreover, the two admissible networks are combinatorially the same. If the admissible network yields a primal feasible solution w.r.t. b, we obtain a pair of primal/dual feasible solutions satisfying complementary slackness and hence the sought optimum solutions x^*, s^*. Suppose the contrary for a contradiction. Then, there is a cut $S \subset V$ s.t. $b(S) > 0$ and there are no ingoing arcs with vanishing reduced costs, i.e., all ingoing arcs have reduced costs of at least 1. Thus, y_v

could be safely increased by 1 for all $v \in S$. Moreover, $\tilde{y}_v$ could be safely increased by $1-\varepsilon$. Thus, the dual objective $\tilde{b}^T\tilde{y}$ would increase by at least $(1-\varepsilon)^2$ because $\mathbb{1}_S^T\tilde{b} = \mathbb{1}_S^T b - \mathbb{1}_S A I_D x \geq 1-\varepsilon$. But this contradicts $\tilde{x}^T\tilde{s} < (1-\varepsilon)^2$. The crossover can be computed in $O(m + n\log n)$ and the transshipment problem in the admissible network that needs to be solved in order to obtain the corresponding primal solution can be solved by a max-flow computation in $O(m^{3/2}\log\frac{n^2}{m}\log U)$ using the algorithm of Goldberg and Rao [14]. □

Note that the additional max-flow computation can be avoided, by using the isolation lemma [23] in order to make the optimal solution unique [9, Lemma 3.12]. However, the perturbation of the cost vector yields non-integral values and, hence, an additional scaling would be necessary in order to make the input integral. Furthermore, this strategy would introduce a new source of randomization.

3 Centering Step

In this section, we describe how to implement the centering step by using a variant of the electrical flow solver from [17]. We would like to stress the fact that we do not rely on the numerical stability analysis of [17] that would involve an additional scaling of the current sources, instead the current sources that are used here can be directly defined as integer values dependent on the current iterates x, s. The centering step takes as input the minor H of G, the variables x, s and the parameter $\mu' = \lceil(1-\tau)\mu\rceil$, with x and s fulfilling $\|\sigma\|_1 \leq \delta$, where $\sigma_a = x_a s_a/\mu - 1$ for all $a \in A_H$. For simplicity, we will denote in this subsection the vectors x and s as vectors over A_H, instead of over A. The goal of `centering_step` is to update x and s to $x+\Delta x, s+\Delta s$ such that $\|\sigma'\|_1 \leq \delta$ holds, where $\sigma'_a = \frac{x'_a s'_a}{\mu'} - 1$. Note that x, s uniquely defines feasible potentials y whose update we denote with Δy. Since we want the new iterates to be interior solutions again, we require

$$A_H\Delta x = 0,\ A_H^T\Delta y + \Delta s = 0 \text{ and } x+\Delta x > 0 \text{ as well as } s+\Delta s > 0. \tag{2}$$

The idea is to compute $\Delta x, \Delta s$ such that $(x+\Delta x)(s+\Delta s) \approx \mu'$, more precisely we will guarantee that $\|\eta\|_1 \leq \delta/4$, where $\eta := \frac{1}{\mu'}[S\Delta x + X\Delta s + Xs] - \mathbb{1}$, i.e., we omit the quadratic term $\Delta x\Delta s$ and allow an approximate solution only. Here $X = \operatorname{diag}(x)$ and $S = \operatorname{diag}(s)$, respectively. It is easy to see that computing such $\Delta x, \Delta s$ satisfying (2) amounts to approximately solving a linear equation system. This approximate solution to the equation system in turn can be found by computing an approximate electrical flow in an electrical network with resistances $r_a = \lceil s_a/x_a\rceil$ and current sources $A_H\varphi^0$, where $\varphi^0_a = x_a - \lceil\mu'/s_a\rfloor$ for all $a \in A_H$. We will denote $R := \operatorname{diag}(r)$. Note that, due to the assumption on the size of β and γ, it holds that $s_a/x_a \geq 1$, see the proof of Lemma 5.

We need some additional notation concerning electrical flows.[6] Let φ be a given flow and let T be a spanning tree of H. For any $a = (v,w) \in A_H\setminus T$, we define $C_a := \{a\}\cup P(w,v)$, where $P(w,v)$ is the unique path in T between w and v, and $r(C_a) := \sum_{a'\in C_a} r_{a'}$ and $\operatorname{tcn}(T) := \sum_{A_H\setminus T}\frac{r(C_a)}{r_a}$ as the *tree condition number of* T. The *tree induced voltages* are defined as $\pi_v := \sum_{a\in P(v_0,v)}\varphi_a r_a$, where v_0 is an arbitrary root of the spanning tree T.

As previously mentioned, our method can be seen as an adaption of the electrical flow solver of Kelner et al. that was introduced in [17]. Kelner et al. showed how to compute an approximate electrical flow φ and corresponding node voltages π in nearly linear time. Up to the choice of the termination criterion and the choice of the initial current our algorithm is identical, see the pseudo-code implementation for further details. In the algorithm, we need to compute a spanning tree and in order to gurantee fast convergence this spanning tree needs to be of low stretch. The current record bounds concerning low stretch spanning trees is achieved by the algorithm of Abraham and Neiman [1]. They show how to compute a tree of stretch $O(m\log n\log\log n)$ in $O(m\log n\log\log n)$

[6] For a more detailed depiction of electrical flows see for example [7, 17].

time.

Algorithm 3: `centering_step`(H, x, s, μ)

Define $r_a := \lceil \frac{s_a}{x_a} \rceil$, $\varphi^0_a := x_a - \lceil \frac{(1-\tau)\mu}{s_a} \rfloor$ for $a \in A_H$.
Let $T :=$spanning tree and $p_a := \frac{r(C_a)}{\mathrm{tcn}(T) r_a} \forall a \in A_H \setminus T$
$\varphi := \varphi^0$, $s' = s$ and $x' = x$
while $\sum_{a \in A_H} |x'_a s'_a - \mu| \geq \delta\mu$ **do**
  Sample $a \in A_H \setminus T$ according to p
  $\alpha := \lceil \frac{-\sum_{a' \in C_a} r_{a'} \varphi_{a'}}{r(C_a)} \rfloor$, $x' \leftarrow x' + \alpha c_a$, $\varphi \leftarrow \varphi + \alpha c_a$
  occasionally: compute tree induced voltages π for φ, $s' := s - A_H^T \pi$
return x', s'

Note that for our purposes here, we will assume the parameter δ to be set to $1/8$ and we will denote with x, s, μ the values of the variables at the beginning of the `centering_step` and the update values we define by $\Delta x := \varphi - \varphi^0$ and $\Delta s := -A_H^T \pi$. We remark that, for our choice of $\delta = 1/8$, evaluating the condition of the while-loop can be done easily with integer arithmetic (by comparing the two integral values resulting from multiplying the inequality by 8).

Moreover, define τ' by $(1-\tau')\mu = \lceil (1-\tau)\mu \rceil = \mu'$ and recall that we set $\varepsilon := \frac{1-\delta}{1+\delta}$. Note that from the minorization and closeness to the central path, we obtain lower and upper bounds on x and s for arcs in the minor, respectively:

$$x_a \in \left[\frac{\varepsilon\beta}{m}, \frac{(1+\delta)\mu m}{\varepsilon\gamma} \right] \quad \text{and} \quad s_a \in \left[\frac{\varepsilon\gamma}{m}, \frac{(1+\delta)\mu m}{\varepsilon\beta} \right] \quad \text{for all } a \in A_H. \tag{3}$$

We first of all show that the updates are feasible moves. The constraints $A_H x' = b$ follow since Δx is a circulation. The update for potentials can be computed in such a way that the dual constraints, except non-negativity of s, are fulfilled as well. For non-negativity of x and s, consider the following lemma - in fact the closeness to the central path already implies non-negativity for x and s, respectively.

Lemma 4. *If $\|\sigma'\|_1 < \delta$, then $x_a + \Delta x_a > 0$ and $s_a + \Delta s_a > 0$ for all $a \in A_H$.*

Proof. If $\|\sigma'\|_1 < \delta$, we must have $|(x_a + \Delta x_a)(s_a + \Delta s_a)/\mu' - 1| < \delta$ for all $a \in A_H$. Thus $(x_a + \Delta x_a)(s_a + \Delta s_a) > 0$ for all $a \in A_H$. Assume $x_a + \Delta x_a \leq 0$ and $s_a + \Delta s_a \leq 0$ for some $a \in A_H$. Then

$$0 \geq s_a(x_a + \Delta x_a) + x_a(s_a + \Delta s_a) = \mu' + s_a x_a > 0, \quad \text{a contradiction.}$$ □

In the following lemma we show that, assuming sufficiently large β and γ, a small $\mathrm{gap}(\varphi, \pi) := \varphi^T R \varphi - 2\pi^T A_H \varphi^0 + \pi^T A_H R^{-1} A_H^T \pi$ of the electrical flow and voltages yields small one-norm of η.

Lemma 5. *Assume $\beta \geq \frac{2^4 m^2}{\delta}$ and $\gamma \geq \frac{2^{17} m^5 \beta U C}{\delta}$. Let $\delta \leq 1/8$ and $\varphi \in \mathbb{R}^{m_H}$ and $\pi \in \mathbb{R}^{n_H}$ such that $\mathrm{gap}(\varphi, \pi) < \frac{2^{-8}\delta^2\mu}{m_H}$. Then it holds that $\|\eta\|_1 \leq \frac{\delta}{4}$, where $\eta := \frac{1}{\mu'}[S\Delta x + X\Delta s + Xs] - \mathbb{1}$.*

Proof. In order to simplify notation, we will use x and s for $M_H x$ and $M_H s$, respectively in this proof, since all consideration are concerned with the minor H. Let us denote with $D_a := \lceil \frac{\mu'}{s_a} \rfloor - \frac{\mu'}{s_a}$

and $E_a := \lceil \frac{s_a}{x_a} \rceil - \frac{s_a}{x_a}$ the errors introduced by the rounding in the algorithm. Then,

$$\begin{aligned}
\|\eta\|_1 &= \sum_{a \in A_H} \frac{x_a}{\mu'} \left| \frac{s_a}{x_a} \Delta x_a + \Delta s_a + s_a - \frac{\mu'}{x_a} \right| \\
&= \sum_{a \in A_H} \frac{x_a}{\mu'} \left| \frac{s_a}{x_a} (\varphi_a - \varphi_a^0) - (\pi_w - \pi_v) + s_a - \frac{\mu'}{x_a} \right| \\
&= \sum_{a \in A_H} \frac{x_a}{\mu'} \left| r_a \varphi_a - (\pi_w - \pi_v) + \frac{s_a}{x_a} D_a - E_a \varphi_a \right| \\
&\leq \frac{\|X(R\varphi - A^T\pi)\|_1}{\mu'} + \frac{\|s\|_1}{2\mu'} + \sum_{a \in A_H} \frac{x_a |\varphi_a|}{\mu'},
\end{aligned} \tag{4}$$

since $|E_a| < 1$ and $|D_a| \leq 1/2$. Moreover, note that using the upper and lower bounds on x and s in (3) and the fact that $\tau' \leq \tau \leq \delta$ yields the following multiplicative errors introduced by the rounding:

$$r_a = \left\lceil \frac{s_a}{x_a} \right\rceil \in \left[\frac{s_a}{x_a}, 2\frac{s_a}{x_a} \right] \quad \text{and} \quad \left\lceil \frac{\mu'}{s_a} \right\rceil \in \left[\frac{\mu'}{2s_a}, \frac{2\mu'}{s_a} \right], \tag{5}$$

where we need that $s_a/x_a \geq 1$, which is guaranteed by $\gamma \geq 2^6 m^3 \beta UC$ and $\beta \geq 4m$. Let us now first show how to bound the last summand in (4): Note that

$$\begin{aligned}
\sum_{a \in A_H} \frac{x_a |\varphi_a|}{\mu'} &\leq \frac{1}{\mu'} \sum_{a \in A_H} x_a \sqrt{\frac{x_a}{s_a}} \cdot \sqrt{\frac{s_a}{x_a}} |\varphi_a| \leq \frac{1}{\mu'} \sqrt{\sum_{a \in A_H} \frac{x_a^3}{s_a}} \sqrt{\sum_{a \in A_H} \frac{s_a}{x_a} |\varphi_a^2|} \\
&\leq \frac{1}{\mu'} \sqrt{\sum_{a \in A_H} \frac{x_a^3}{s_a}} \|\varphi\|_R.
\end{aligned}$$

Since the energy of the initial flow φ^0 is at least the energy of φ, we obtain

$$\begin{aligned}
\|\varphi\|_R^2 \leq \|\varphi^0\|_R^2 &\leq \sum_{a \in A_H} \frac{2s_a}{x_a} \left[x_a^2 - 2x_a \left\lceil \frac{\mu'}{s_a} \right\rceil + \left\lceil \frac{\mu'}{s_a} \right\rceil^2 \right] \leq \sum_{a \in A_H} 2x_a s_a - 2\mu' + \frac{8\mu'^2}{x_a s_a} \\
&\leq m_H \left[2(1+\delta)\mu - 2(1-\tau')\mu + \frac{8\mu(1-\tau')^2}{(1-\delta)} \right] \leq \frac{1}{2}\mu + \frac{8\mu(1-\tau')^2}{(1-\delta)} \leq 10\mu m_H,
\end{aligned} \tag{6}$$

where we used $x_a s_a/\mu \geq 1-\delta$, $\delta \leq \frac{1}{8}$ and $\tau' \leq \tau = \delta/\sqrt{m} \leq \delta$. The upper bound for x and the lower bound for s in (4) yield $\sqrt{\sum_{a \in A_H} x_a^3/s_a}/\mu' \leq \sqrt{(1+\delta)\mu m^5}/(\varepsilon^3\gamma^2)$, again using $\tau' \leq \delta$. Putting the two bounds together, we obtain $\sqrt{10(1+\delta)}\mu m^3/(\varepsilon^3\gamma^2)$ as a bound on the third summand in (4). Now, for the first summand in (4), note that $\text{gap}(\varphi, \pi) = \sum_{a \in A_H \setminus T} [r_a \varphi_a - (\pi_w - \pi_v)]^2 / r_a$, as shown in [17, Lemma 4.4]. Hence,

$$\frac{\|X(R\varphi - A^T\pi)\|_1}{\mu'} = \sum_{a \in A_H \setminus T} \frac{|r_a \varphi_a - (\pi_w - \pi_v)|}{\sqrt{r_a}} \frac{\sqrt{r_a} x_a}{\mu'} \leq \frac{\sqrt{\text{gap}\, \|X^2 r\|_1}}{\mu'} \leq \frac{2\sqrt{\text{gap}\, m_H}}{\sqrt{\mu}}$$

using the Cauchy-Schwarz-Inequality, $x_a s_a/\mu \leq 1+\delta$ and again $\tau' \leq \delta$. Finally, for the second summand in (4), note that $\frac{\|s\|_1}{2\mu'} \leq \frac{(1+\delta)\mu m^2}{2\mu' \varepsilon \beta} \leq \frac{m^2}{2\varepsilon^2 \beta}$. Hence, plugging all bounds together into (4) yields

$$\|\eta\|_1 \leq \frac{2\sqrt{\text{gap}\, m}}{\sqrt{\mu}} + \frac{m^2}{2\varepsilon^2\beta} + \frac{\sqrt{10(1+\delta)}\mu m^3}{\varepsilon^3\gamma^2} \leq \frac{\delta}{8} + \frac{\delta}{16} + \frac{\delta}{16} \leq \frac{\delta}{4},$$

with $\beta \geq \frac{2^4 m^2}{\delta}$, $\gamma \geq \frac{2^{12} m^4 \beta UC}{\delta}$, $\text{gap} \leq \frac{2^{-8}\delta^2 \mu}{m_H}$ and the upper bound on $\mu \leq \mu_0 \leq 24 m \beta \gamma UC$. □

We next show that a small one-norm of η actually implies the closeness to the central path of the new iterates x' and s', again assuming sufficiently large β and γ.

Lemma 6. *Let $\delta \leq 1/8$. If $\|\eta\|_1 \leq \frac{\delta}{4}$, then $\|\sigma'\|_1 \leq \delta$ where $\sigma'_a := \frac{x'_a s'_a}{\mu'} - 1$ for $a \in A_H$.*

Proof. By definition of η, it holds that $(1+\eta_a)\mu' = \Delta x_a s_a + x_a \Delta s_a + x_a s_a$, which yields

$$\Delta x_a \Delta s_a = \tfrac{1}{2x_a s_a}\left[((1+\eta_a)\mu' - x_a s_a)^2 - (\Delta x_a s_a)^2 - (x_a \Delta s_a)^2\right] \quad \text{for } a \in A_H. \tag{7}$$

For the one-norm of σ', we get the following estimate

$$\|\sigma'\|_1 = \sum_{a\in A_H}\left|\frac{x'_a s'_a}{\mu'} - 1\right| = \sum_{a\in A_H}\left|\frac{\Delta x_a \Delta s_a}{\mu'} + \eta_a\right| \leq \sum_{a\in A_H}\left|\frac{\Delta x_a \Delta s_a}{\mu'}\right| + \|\eta\|_1.$$

Using (7) and the triangle inequality, we obtain for the first summand that

$$\begin{aligned}\sum_{a\in A_H}\left|\frac{\Delta x_a \Delta s_a}{\mu'}\right| &\leq \sum_{a\in A_H}\frac{[(1+\eta_a)\mu' - x_a s_a]^2}{2x_a s_a \mu'} + \frac{(\Delta x_a s_a)^2 + (x_a \Delta s_a)^2}{2x_a s_a \mu'} \\ &= \sum_{a\in A_H}\frac{[(1+\eta_a)\mu' - x_a s_a]^2}{x_a s_a \mu'},\end{aligned}$$

because Δx and Δs are orthogonal and thus the sum over all arcs in (7) yields the last equality. Now let τ' be such that $\mu' = (1-\tau')\mu$. Factoring out μ, plugging in the definition of μ' and σ_a and using the lower bound on $\frac{x_a s_a}{\mu}$ yields

$$\sum_{a\in A_H}\left|\frac{\Delta x_a \Delta s_a}{\mu'}\right| \leq \sum_{a\in A_H}\frac{\mu^2\left(\frac{(1+\eta_a)\mu'}{\mu} - \frac{x_a s_a}{\mu}\right)^2}{x_a s_a \mu'} \leq \sum_{a\in A_H}\frac{[(1-\tau')(1+\eta_a) - 1 - \sigma_a]^2}{(1-\tau')(1-\delta)}.$$

Expanding the square, re-grouping the terms and using $\mathbb{1}^T\sigma \leq \|\sigma\|_1 \leq \delta$, $-\mathbb{1}^T\eta \leq \|\eta\|_1$, the equivalence of norms and the Cauchy-Schwarz-Inequality for $-\eta^T\sigma \leq \|\eta\|_2\|\sigma\|_2$, yields

$$\sum_{a\in A_H}\left|\frac{\Delta x_a \Delta s_a}{\mu'}\right| \leq \frac{(1-\tau')^2\|\eta\|_1^2 + [2(1-\tau')\tau' + 2\delta(1-\tau')]\|\eta\|_1 + \delta^2 + 2\delta\tau' + \tau'^2 m}{(1-\tau')(1-\delta)}.$$

Using $\tau' \leq \tau = \frac{\delta}{\sqrt{m}} \leq \delta$ and $1-\tau' \leq 1$ yields

$$\|\sigma'\|_1 \leq \frac{\|\eta\|_1^2 + 4\delta\|\eta\|_1 + 4\delta^2 + (1-\tau')(1-\delta)\|\eta\|_1}{(1-\tau')(1-\delta)} \leq \frac{\|\eta\|_1^2 + (4\delta+1)\|\eta\|_1 + 4\delta^2}{(1-\delta)^2}$$

and $\|\eta\|_1 \leq \frac{\delta}{4}$ gives the result for any $\delta \leq \frac{1}{8}$. □

It remains to argue the convergence behavior of `centering_step`. The following lemma can be seen as the equivalence of [17, Lemma 4.5]. We obtain an additional factor of 3/4 due to the rounding to the nearest integer of α. We can show that assuming large enough β and γ, the rounding does not impede the convergence speed significantly.

Lemma 7. *The expected decrease in primal energy $\|\varphi\|_R^2$ for every cycle-update in the inner while-loop of Algorithm 1 is at least $\frac{3\,\text{gap}}{4\,\text{tcn}(T)}$, provided that $\beta \geq 2^5 m^3/\delta$, $\gamma \geq 2^6 m^3 \beta UC$ and $\text{gap} \geq 2^{-8}\delta^2\mu/m_H$.*

Proof. Let $\Lambda_a := \sum_{a'\in C_a} r_{a'}\varphi_{a'} = r_a\varphi_a - (\pi_w - \pi_v)$. In every iteration of the inner while-loop an arc a is sampled and φ is updated to $\varphi' = \varphi - \alpha_a\theta_a$, denoting $\alpha_a := \lceil \frac{\Lambda_a}{r(C_a)} \rfloor$ and θ_a the indicator vector corresponding to C_a. The change in energy due to this update is $\|\varphi - \alpha_a\theta_a\|_R^2 - \|\varphi\|_R^2 =$

$-2\alpha_a\Lambda_a + \alpha^2 r(C_a)$. Thus, using the definition of the probabilities $p_a = \frac{r(C_a)}{\operatorname{tcn}(T)r_a}$ for $a \in A_H \setminus T$, the expected decrease in energy is

$$\mathrm{E}\left[\|\varphi'\|_R^2 - \|\varphi\|_R^2\right] = \sum_{A_H \setminus T} \frac{r(C_a)^2\alpha_a^2}{\operatorname{tcn}(T)r_a} - \frac{2r(C_a)\alpha_a\Lambda_a}{\operatorname{tcn}(T)r_a} = \sum_{A_H \setminus T} \frac{\left[\Lambda_a - \alpha_a r(C_a)\right]^2 - {\Lambda_a}^2}{\operatorname{tcn}(T)r_a}.$$

Now, using $\operatorname{gap} = \sum_{A_H \setminus T} {\Lambda_a}^2/r_a$ and $r(C_a) \le \|r\|_1$ for all $a \in A_H \setminus T$ for the right term in the enumerator and $|\Lambda_a - \alpha_a r(C_a)| \le r(C_a)/2$ as the rounding error for the left term, yields

$$\mathrm{E}\left[\|\varphi'\|_R^2 - \|\varphi\|_R^2\right] \le \frac{1}{4\operatorname{tcn}(T)}\|r\|_1 \operatorname{tcn}(T) - \frac{\operatorname{gap}}{\operatorname{tcn}(T)} \le \frac{1}{4\operatorname{tcn}(T)}\|r\|_1 m^2 - \frac{\operatorname{gap}}{\operatorname{tcn}(T)}.$$

Using $r_a \le 2s_a/x_a$, which follows from $\gamma \ge 2^6 m^3 \beta UC$, see (5) and the upper and lower bounds on x, s in (3), yields

$$\|r\|_1 m^2 \le \frac{2(1+\delta)m^5\mu}{\varepsilon^2\beta^2} \le \frac{2(1+\delta)m^5\mu}{\varepsilon^2}\frac{\delta^2}{2^{10}m^6} \le \frac{(1+\delta)}{2\varepsilon^2}\frac{\delta^2\mu}{2^8 m} \le \frac{\delta^2\mu}{2^8 m} \le \operatorname{gap},$$

where we used $\beta \ge 2^5 m^3/\delta$. This yields $\mathrm{E}\left[\|\varphi'\|_R^2 - \|\varphi\|_R^2\right] \le \frac{-3\operatorname{gap}}{4\operatorname{tcn}(T)}$. □

The above lemma yields a corresponding result to [17, Theorem 3.2], that is an expected guarantee on the approximation of the primal and dual energy. However, as Kelner et al. state "one can use Markov bound or Chernoff bound to prove a probabilistic but exact guarantee". In fact, the near-linear run-time of the `centering_step` follows as in the proof of Theorem 2 in [6].

4 Conclusion

We conclude by deducing the precise bound on the size of the numbers that the algorithm has to deal with: Note that the voltages, the values of the α, as well as the currents φ are each bounded in absolute value by $\|R\varphi\|_1$. By using the Cauchy-Schwarz-Inequality, we obtain $\|R\varphi\|_1 \le \sqrt{\|r\|_1} \cdot \|\varphi\|_R$. The first factor can be bounded by $\sqrt{2(1+\delta)\mu_0 m^3/\varepsilon\beta}$ using the bounds from (5) and (3). The second, analogously to the estimation in (6), can be bounded by $\sqrt{10\mu m}$. Together with $\mu \le \mu_0 = 3m\beta\gamma UC/\delta$, which is shown in Appendix A and using $\delta = 1/8$, we obtain $2^8 m^3 \gamma UC$ as the final bound on all numbers appearing in the centering step. The values of x and s are bounded by $(1+\delta)\mu m/(\varepsilon\gamma)$ and $(1+\delta)\mu m/(\varepsilon\beta)$, respectively, which yields the bounds of $2^6 m^2 \beta UC$ and $2^6 m^2 \gamma UC$. This shows Theorem 2. Finally, an arbitary given min-cost flow instance can first be scaled down, i.e., divided by its gcd, and then scaled up again with β and γ of size as described in Theorem 2. It follows that we can always guarantee the bound on the size of numbers in Theorem 1.

References

[1] Ittai Abraham and Ofer Neiman. Using Petal-Decompositions to Build a Low Stretch Spanning Tree. In Howard J. Karloff and Toniann Pitassi, editors, *STOC*, pages 395–406. ACM, 2012.

[2] Ravindra K. Ahuja, Andrew V. Goldberg, James B. Orlin, and Robert Endre Tarjan. Finding minimum-cost flows by double scaling. *Math. Program.*, 53:243–266, 1992.

[3] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows - Theory, Algorithms and Applications*. Prentice Hall, 1993.

[4] Ernst Althaus and Kurt Mehlhorn. Maximum network flow with floating point arithmetic. *Inf. Process. Lett.*, 66(3):109–113, 1998.

[5] Ruben Becker, Maximilian Fickert, and Andreas Karrenbauer. A novel dual ascent algorithm for solving the min-cost flow problem. In Michael T. Goodrich and Michael Mitzenmacher, editors, *Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments, ALENEX 2016, Arlington, Virginia, USA, January 10, 2016*, pages 151–159. SIAM, 2016.

[6] Ruben Becker and Andreas Karrenbauer. A Simple Efficient Interior Point Method for Min-Cost Flow. In Hee-Kap Ahn and Chan-Su Shin, editors, *Algorithms and Computation - 25th International Symposium, ISAAC 2014, Jeonju, Korea, December 15-17, 2014, Proceedings*, volume 8889 of *Lecture Notes in Computer Science*, pages 753–765. Springer, 2014.

[7] B. Bollobas. *Modern Graph Theory*. Graduate Texts in Mathematics. Springer New York, 1998.

[8] Michael B. Cohen, Aleksander Madry, Piotr Sankowski, and Adrian Vladu. Negative-weight shortest paths and unit capacity minimum cost flow in $\tilde{o}(m^{10/7} \log W)$ time. *CoRR*, abs/1605.01717, 2016.

[9] Samuel I. Daitch and Daniel A. Spielman. Faster Approximate Lossy Generalized Flow via Interior Point Algorithms. In Cynthia Dwork, editor, *STOC*, pages 451–460. ACM, 2008. Full version available under http://arxiv.org/abs/0803.0988.

[10] Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. In Michael Jünger, Gerhard Reinelt, and Giovanni Rinaldi, editors, *Combinatorial Optimization - Eureka, You Shrink!, Papers Dedicated to Jack Edmonds, 5th International Workshop, Aussois, France, March 5-9, 2001, Revised Papers*, volume 2570 of *Lecture Notes in Computer Science*, pages 31–33. Springer, 2001.

[11] D. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, NJ, USA, 2010.

[12] Andrew V. Goldberg. An efficient implementation of a scaling minimum-cost flow algorithm. *J. Algorithms*, 22(1):1–29, 1997.

[13] Andrew V. Goldberg and Michael Kharitonov. On implementing scaling push-relabel algorithms for the minimum-cost flow problem. In David S. Johnson and Catherine C. McGeoch, editors, *Network Flows And Matching, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, October 14-16, 1991*, volume 12 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 157–198. DIMACS/AMS, 1991.

[14] Andrew V. Goldberg and Satish Rao. Beyond the Flow Decomposition Barrier. In *FOCS*, pages 2–11, 1997.

[15] Andrew V. Goldberg and Robert E. Tarjan. Finding minimum-cost circulations by successive approximation. *Math. Oper. Res.*, 15(3):430–466, July 1990.

[16] Andrew V. Goldberg and Robert Endre Tarjan. Solving minimum-cost flow problems by successive approximation. In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 7–18. ACM, 1987.

[17] Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A Simple, Combinatorial Algorithm for Solving SDD Systems in Nearly-Linear Time. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *STOC*, pages 911–920. ACM, 2013.

[18] Lutz Kettner, Kurt Mehlhorn, Sylvain Pion, Stefan Schirra, and Chee-Keng Yap. Classroom examples of robustness problems in geometric computations. *Comput. Geom.*, 40(1):61–78, 2008.

[19] Zoltán Király and P. Kovács. Efficient implementations of minimum-cost flow algorithms. *CoRR*, abs/1207.6381, 2012.

[20] Ioannis Koutis, Gary L. Miller, and Richard Peng. Approaching Optimality for Solving SDD Linear Systems. In *FOCS*, pages 235–244. IEEE Computer Society, 2010.

[21] Yin Tat Lee, Richard Peng, and Daniel A. Spielman. Sparsified cholesky solvers for SDD linear systems. *CoRR*, abs/1506.08204, 2015.

[22] Yin Tat Lee and Aaron Sidford. Path Finding Methods for Linear Programming: Solving Linear Programs in $\tilde{O}(\sqrt{rank})$ Iterations and Faster Algorithms for Maximum Flow. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 424–433. IEEE Computer Society, 2014.

[23] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.

[24] James B. Orlin. A Faster Strongly Polynominal Minimum Cost Flow Algorithm. In Janos Simon, editor, *STOC*, pages 377–387. ACM, 1988.

[25] J.B. Orlin. *Genuinely Polynomial Simplex and Non-simplex Algorithms for the Minimum Cost Flow Problem*. Report OS / Centrum voor Wiskunde en Informatica. Stichting Mathematisch Centrum, 1984.

[26] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In László Babai, editor, *STOC*, pages 81–90. ACM, 2004.

[27] Éva Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5(3):247–256, 1985.

[28] Pravin M. Vaidya. Speeding-up linear programming using fast matrix multiplication (extended abstract). In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 332–337. IEEE Computer Society, 1989.

[29] Yinyu Ye. *Interior point algorithms - theory and analysis*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1998.

A Appendix: Finding Arbitrarily Central Points

In this section, we show how to construct an auxiliary instance and corresponding feasible interior solutions that are arbitrarily close to the central path. In fact, we can use the same technique as used in [6] for a potential reduction method. Let us denote with (G^0, b^0, u, c^0) the given input instance. Let V^0 denote the node set of G^0 and let A^0 denote its arc set. The method constructs an auxiliary instance (G, b, c) that is an un-capacitated min-cost flow problem and corresponding primal and dual interior solutions for it. We denote with V and A the node and arc set of G, respectively, and with $n = |V|$ and $m = |A|$ their respective cardinalities. The first part of the construction is illustrated in Figure 1 from left to middle. For each $a = (v, w) \in A^0$, insert the two nodes v, w into G as well as a new node vw together with arcs $\acute{a} = (v, vw)$ and $\grave{a} = (w, vw)$, read "*a up*" and "*a down*". Define the new costs by $c_{\acute{a}} := c_a$ and $c_{\grave{a}} := 0$. The new demand vector is given by $b_v := b^0_v - u(\delta^{\text{in}}(v))$ for nodes $v \in V^0$ and $b_{vw} := u_{(v,w)}$ for the newly inserted nodes. A flow of f on $a \in G^0$ corresponds to a flow of f on $\acute{a} = (v, vw)$ and a flow of $u_a - f$ on $\grave{a} = (w, vw)$ in G. This is a well-known construction and it is known that the new instance is equivalent to the old instance, see for example [3, Section 2.4].

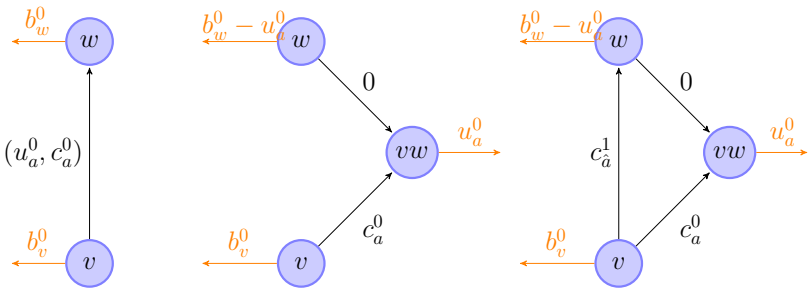

Figure 1: The transition from the left to the middle, which is done for each arc, removes the capacity constraint. From the middle to the right: In order to balance the $x_a s_a$, we introduce the arc $\hat{a} = (v, w)$ with high cost $c_{\hat{a}}$ and reroute flow along it. The direction of $\hat{a}$ depends on a tree solution z in the original graph. It is flipped if $z_a \leq u_a/2$. All arcs in the middle and the right graph have infinite capacity. Only the costs are shown.

The goal is to find values x, y, s that are primal and dual feasible, integral, as well as close to the central path. The latter means that we want that $\sum_{a \in A} |\frac{x_a s_a}{\mu_0} - 1| \leq \delta$ for some parameters $\mu_0 > 0$ and $\delta \geq 0$. We show that the technique from [6] can actually achieve this. With a different choice of parameters, the technique yields closeness to the central path for arbitrary δ.

The second part of the construction starts by computing an integral (not necessarily feasible) tree solution z in G^0 for an arbitrary spanning tree T, set $x_{\grave{a}} = x_{\acute{a}} = u_a/2$ and introduce the additional arc $\hat{a} = \text{sign}(z_a - u_a/2)(v, w)$. If $z_a - u_a/2 = 0$, do not introduce the arc $\hat{a}$. Set the flow on $\hat{a}$ to $x_{\hat{a}} = |z_a - u_a/2|$ and the cost to $c_{\hat{a}} = \lceil t/|z_a - u_a/2| \rceil$. A sufficiently large choice of t guarantees that no flow is routed across $\hat{a}$ in an optimal solution. More precisely, for $t \geq 2\beta\gamma mUC$, we obtain $c_{\hat{a}} \geq t/(2\beta U) \geq m\gamma C \geq \mathbb{1}^T c^0$, where we used that $z_a \leq \|b^0\|_1/2$. It follows that the cost of $\hat{a}$ is at least the cost of any path in G^0 and thus introducing $\hat{a}$ does not change the optimum of the problem. The dual variables are set as $y_v, y_w = 0$ and $y_{vw} = -2t/u_a$, which determines $s_{\acute{a}} = c_a + 2t/u_a$, $s_{\grave{a}} = 2t/u_a$ and $s_{\hat{a}} = c_{\hat{a}} = \lceil t/|z_a - u_a/2| \rceil$. We can now show that x, y, s are arbitrarily close to the central path.

Lemma 8. *Let x, y, s as described above.*

1. *For any $t > 0$, it holds that $x_a s_a \in [t, t + \beta\gamma UC]$ for all $a \in A$.*
2. *Let $t \geq \beta\gamma UC(\frac{m}{\delta} - 1)$ and $\mu_0 := t + \beta\gamma UC$, then $\|\sigma\|_1 \leq \delta$, where $\sigma_a = \frac{x_a s_a}{\mu_0} - 1$.*

Proof. 1. Let $a \in A_0$ be any arc in G_0, then $x_{\acute{a}} = x_{\grave{a}} = u_a/2$. It follows that,

$$x_{\acute{a}} s_{\acute{a}} = \frac{u_a}{2}\left(c_{\acute{a}} + \frac{2t}{u_a}\right) = t + \frac{u_a c_a}{2} \leq t + \frac{\beta\gamma UC}{2}, \quad x_{\grave{a}} s_{\grave{a}} = \frac{u_a}{2}\left(c_{\grave{a}} + \frac{2t}{u_a}\right) = t \quad \text{and}$$

$$x_{\hat{a}} s_{\hat{a}} \geq \left|z_a - \frac{u_a}{2}\right| \frac{t}{|z_a - \frac{u_a}{2}|} = t \quad \text{and} \quad x_{\hat{a}} s_{\hat{a}} \leq \left|z_a - \frac{u_a}{2}\right| \left(\frac{t}{|z_a - \frac{u_a}{2}|} + 1\right) \leq t + \beta U.$$

2. From the first part, it follows that $x_a s_a \in [t, t + \beta\gamma UC]$ for $a \in A$ or equivalently $\frac{x_a s_a}{\mu_0} \in [t/(t + \beta\gamma UC), 1] = [1 - \beta\gamma UC/(t + \beta\gamma UC), 1]$. This yields $\frac{x_a s_a}{\mu_0} \geq 1 - \delta/m$ with the lower bound on t and hence

$$\|\sigma\|_1 = \sum_{a \in A} \left|\frac{x_a s_a}{\mu_0} - 1\right| \leq \sum_{a \in A} \frac{\delta}{m} = \delta. \qquad \square$$

From the above lemma and the lower bound on t that ensures that the optimum is not changed due to the introduction of the new arcs $\hat{a}$, we obtain that choosing $\mu_0 = 3m_0\beta\gamma UC/\delta$ is sufficient. Choosing $\delta = 1/8$ implies that setting $\mu_0 = 24m\beta\gamma UC$ is sufficient.